A Distributed, Scalable, Community Care Network Architecture for Wide-Area Electronic Patient Records: Modeling and Simulation

S. Ghosh & K. Han, Department of Computer Science and Engineering
Arizona State University, Tempe, AZ 85287
&
R. Reddy, S. Reddy, S. Kankanahalli, J. Jagannathan, & R. Shank
CERC, West Virginia University, Morgantown, WV 26506

Abstract

Principal systems issues relative to computerizing patient medical records that are yet to be addressed in the scientific literature include (1) the characteristics of networks, i.e. bandwidth and capacity, and their impact on the performance of the system, (2) the architecture and the underlying algorithm of the system, (3) the location and migration of medical records, (4) scalability of the system, and (5) the nature of the performance variation under heavy and light use of the network. Key parameters that affect performance include the number of patients, doctors, frequency of patient visits, and the number of electronic queries and record entries initiated during a patient-doctor interaction episode.

This paper presents AMPReD, a Distributed, Scalable, Community Care Network Architecture that aims to provide Real-Time Access to Geographically-Dispersed Patient Medical Records. The AMPReD model includes stationary hospitals and medical clinics, mobile clinics, migrating doctors as well as patients, the communications network, and the patient medical record database. AMPReD's goals include (1) the accurate modeling of the propagation of medical records and (2) providing real-time access to patient medical records from anywhere in the system. To achieve these goals, an asynchronous, distributed algorithm must be developed that achieves concurrent access of multiple, autonomous databases. AMPReD is modeled and simulated for a representative community care network on a network of workstations configured as a loosely-coupled parallel processor, for different parametric combinations of number of doctors, patients, and number of queries or record entries generated corresponding to every patient-doctor interaction episode. AMPReD defines and obtains key performance measures including the idle times of the doctors, patient waiting times, the access times of queries as functions of their sizes, and the growth of the databases. In addition, AMPReD also measures the deviation of the actual time required for a patient-doctor interaction episode from the scheduled interaction interval, as a function of the network load. For the representative system selected, performance measures indicate that the network, utilizing $\frac{1}{2}T_1$ links, and the database system poses no bottleneck to the system even where the number of doctors and patients within a 30 minute interval are chosen at 192 and 200 respectively. A T_1 is a standard, digital, transmission link that is rated at 1.44Mbits/sec.

1. Introduction

Community care networks that aim to integrate geographically-dispersed patient medical records promise to enable real-time access of patient records to the medical community during critical times, reduce duplication of tests, achieve better efficiency of resource usage, lower medical costs, and improve the quality of medical decision-making through greater availability of accurate information. The modeling and simulation of such large-scale networks on the computer prior to building a prototype, is necessitated by economic efficiency and the desire to develop a system with characteristics close to its original specifications. The scientific literature reports only a handful of studies related to the modeling and simulation of such networks under representative traffic scenarios. McDaniel [1] reports simulation studies of a wide area health care network in Canada. The study is limited in that it utilizes 2400 baud telephone modem connection, the data message size ranges from 175 to 2000 bytes, and the frequency of records transferred is low. It reports that 80,000 messages are exchanged between 1553 doctors, 26 hospitals, four medical labs, one provincial lab, and one insurer but fails to report the number of patients served in the simulation. In sharp contrast to the US culture, in the Canadian medical custom, the records appear to be oriented around doctor, insurance provider and government. Evidently, this reflects the sharp contrast in the architectural design of the US system, relative to the Canadian system. Significant weaknesses in [1] include the total absence of how data is routed and the failure to model the transmission of radiographic images. Since the simulation timestep is unreported, the accuracy of the simulation is uncertain.

2. A Distributed Community Care Network Architecture, AMPReD

To achieve AMPReD's goals, this paper introduces an asynchronous, distributed algorithm design that achieves concurrent access of multiple, autonomous databases. The algorithm is expected to deliver scalability, i.e. it should service increasing numbers of patients accompanied by increasing numbers of doctors, medical vans, hospitals and medical clinics, while providing the same level of performance. AMPReD is modeled and simulated for a representative community care network on a network of workstations configured as a loosely-coupled parallel processor, for different parametric combinations of number of doctors, patients, and number of queries or record entries generated corresponding to every patient-doctor interaction episode.

In AMPReD, each hospital and clinic operates as before, makings its own decisions autonomously, but with AMPReD executing in the background. While the doctors are associated with a hospital or clinic, they do migrate, stochastically, to other clinics and hospitals. Patients are associated with a primary medical center but they too are free to migrate to different hospitals or clinics. Furthermore, entered data is patient oriented, i.e. while the data is stored in the institution or hospital or with the doctor, the data can and is released to the patients. In general, summaries of a patient's medical record are stored at the primary medical center. Where the primary medical centers are grouped as clinics, the information is duplicated at the group's headquarters (HQ). Data is not stored at a central facility, instead, patient data is distributed over the network to increase reliability, robustness, efficiency, and resistant to catastrophic failure. AMPReD assumes that the database possesses infinite storage capacity and argues that the access time of data is independent of the current size of the database. The difference between hospitals and clinics are as follows. While a hospital's information is self-contained in a database, clinics report to their HQs. Also, a hospital maintains one database while a clinic system may have multiple databases.

In the AMPReD network, all links are assumed to be $\frac{1}{2}T_1$ (telephone trunk) lines (0.772 Mbits/sec) in each direction and the interconnection network resembles a partially connected network. In a partially interconnected network, not all nodes are connected directly to each other. As a result, a message from a node to a second node may need to propagate via a third node. Evidently, this underscores the need for routing messages through the network. In AMPReD, routing is virtual path oriented to improve performance. Every node is aware of the static topology of the entire network. As a result of virtual path routing, superior performance is expected. For every patient-doctor interaction episode, AMPReD assumes that the pattern of a patient's visit to the doctor is pseudo-random with Gaussean time intervals between the visits. Figure 1 shows a pictorial representation of the representative community care network with 2 hospitals A and B and two clinic systems 1 and 2. Each clinic system maintains a headquarter and 6 clinics.

3. Simulation Results and Performance Analysis

The distribution of patient assertion across the different hospitals or medical centers of the system and the nature of the interactions are generated through the use of pseudo-random number generators.

The representative medical system shown in Figure 1, consists of 16 medical centers or hospitals each of which is linked with high-speed line, rated at $\frac{1}{2}T_1$. Patients are asserted pseudo-randomly at different medical centers in the system and the frequency with which patients migrate among the medical centers is assumed given by a

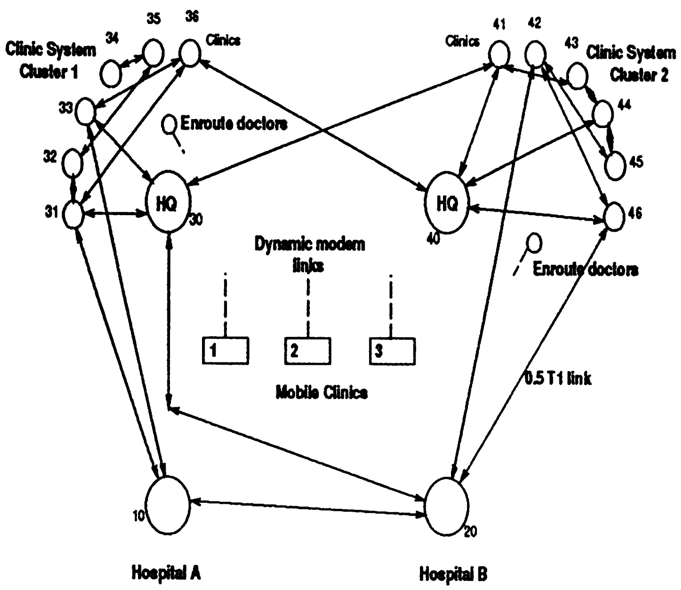

Figure 1: Community Care Network Architecture

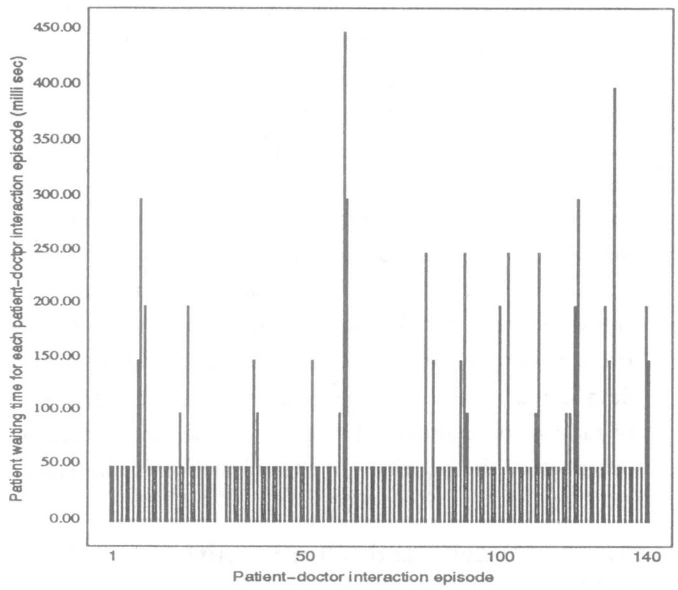

Figure 2: Patient Waiting Time

normal distribution with $\mu = 1$ and $\sigma = 1$. For accuracy in the simulation of the network, the timestep is set at 50ms. A timestep is the basic unit of time in the simulation such that two consecutive events are separated by one or more timesteps. The fine resolution of the timestep, i.e. 50ms, guarantees very accurate simulation at the packet level and it is very time consuming. Each simulation run corresponds to 30 minutes of actual operational time, executes on 16 SUN sparc 10 workstations, and requires approx. 3 hours of wall clock time. The size of the queries are generated stochastically and they range from 500 bytes for a simple medical record to 2 Mbytes for an X-ray image. A number of simulation experiments are conducted. First, the number of doctors is set at 144 while the corresponding number of patients is 100. Second, the number of doctors is set at 192 while the corresponding number of patients is 200. The behavior of the graphs for the second set of results parallel those for the first experiment and are not presented here.

Figure 2 presents the distribution of the patient waiting time as a function of the patient-doctor interaction episode. The average is computed as 0.079 second while the maximum waiting time is recorded as 0.45 second.

Figure 3 presents the variation of the response time required by the queries as a function of the packet size. Clearly, as the packet size increases, the network experi-

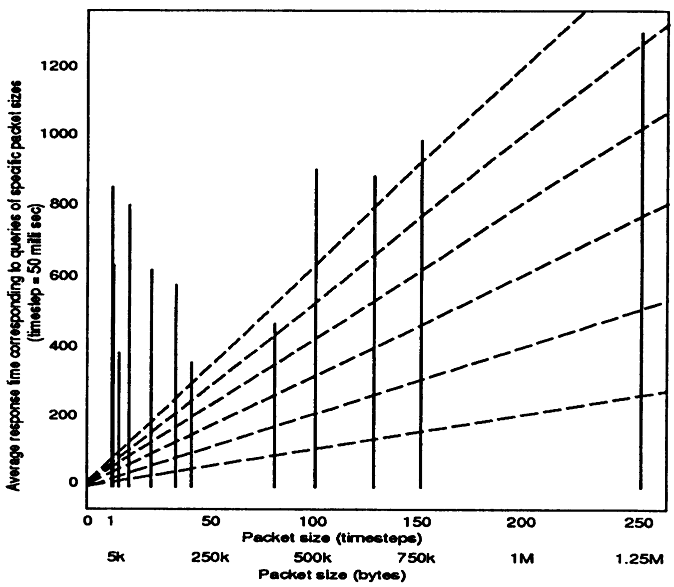

Figure 3: Average Response Time for Queries of Specific Packet Sizes

ences greater load and the response time increases.

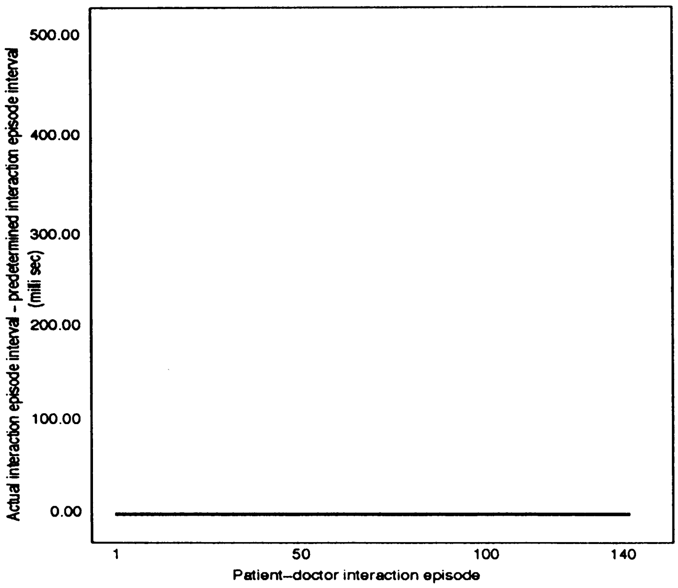

Figure 4: Deviation of Actual Time required for an Interaction Episode Relative to the Originally Scheduled Time Interval, as a function of the Patient-Doctor Interaction Episodes

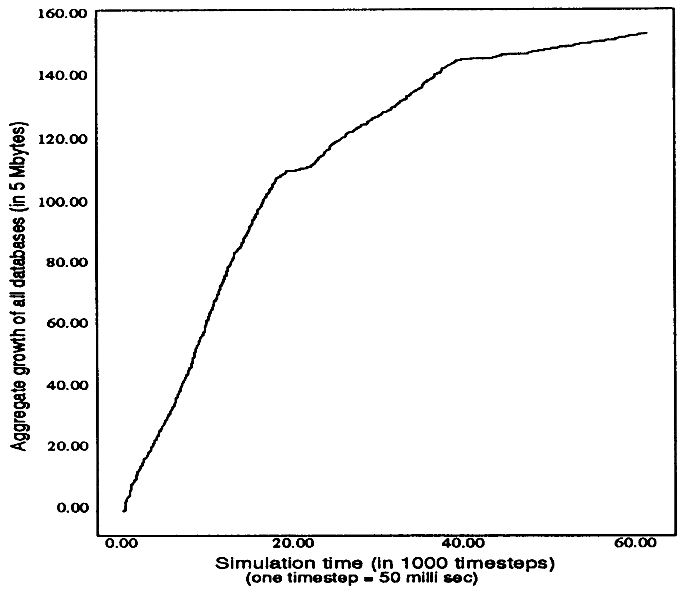

Figure 5: Aggregate Growth of all Databases

Figure 4 reveals that the network loading does not delay the completion of a patient-doctor interaction episode, relative to the originally scheduled time interval. Thus, an interaction episode is completed within the planned interval, attesting to the network's superior performance.

Figure 5 shows the aggregate growth of all databases within the length of the simulation that corresponds to 30 minutes of actual operation. The graph attests to the significant size of data being exchanged through the network and is useful towards planning the long-term behavior of the system.

Figure 6 presents the doctors free times as a function of the patient-doctor interaction episode. While the average free time is 2.9 seconds, the maximum free time for a doctor is observed to be 36.6 minutes. Clearly, the doctor who is free for 36.6 minutes never interacts with any patient during the course of this simulation and it is a result of the assumption that there are more doctors (144) relative to the number of patients (100).

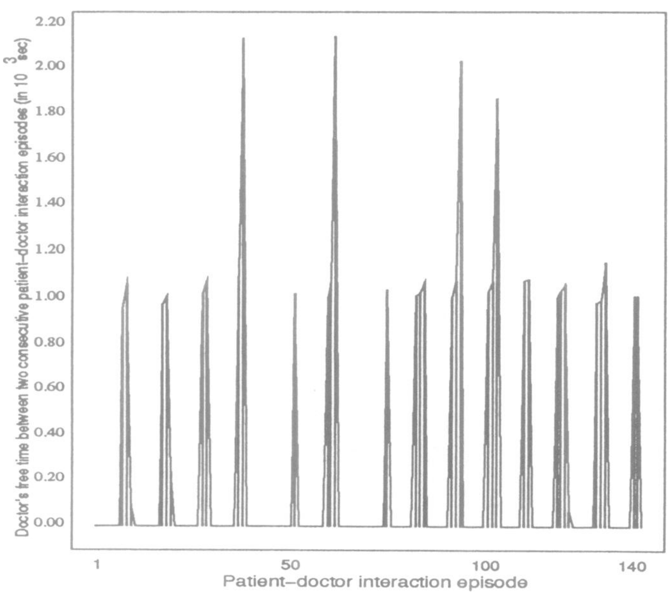

Figure 6: Doctor's Free Time

Acknowledgements

The authors gratefully acknowledge the support of the National Library of Medicine through contract N01-LM-3-3525. The authors also thank Tony Lee at Brown University for his valuable assistance.

References

[1] James G. McDaniel, "Simulation Studies of a Wide Area Health Care Network," Proceedings of the 18th Annual Symposium on Computer Applications in Medical Care (SCAMC), Nov 5-9, 1994, pp. 438-444.

[2] Donald J. Hillman, "A Simulation model of the ELHILL3 system of the National Library of Medicine," ONLINE REVIEW, Vol. 4, No. 3, 1980.

[3] Scott Wallace, "The Computerized Patient Record," BYTE, May 1994, pp. 67-75.

[4] Averill M. Law and W. David Kelton, Simulation Modeling and Analysis, McGraw Hill Inc., Mew York, 1991.